Submitted to the Annals of Applied Statistics
arXiv: stat.ME/0000.0000

arXiv:1403.6168v1 [stat.ME] 24 Mar 2014

STRUCTURED REGULARIZATION FOR CONDITIONAL GAUSSIAN GRAPHICAL MODELS

By Julien Chiquet, Tristan Mary-Huard, Stéphane Robin.

UMR 8071 CNRS/UEVE and UMR 518 INRA/AgroParisTech

We propose a regularized method for multivariate linear regression when the number of predictors may exceed the sample size. This method is designed to strengthen the estimation and the selection of the relevant input features with three ingredients: it takes advantage of the dependency pattern between the responses by estimating the residual covariance; it performs selection on direct links between predictors and responses; and selection is driven by prior structural information. To this end, we build on a recent reformulation of the multivariate linear regression model to a conditional Gaussian graphical model and propose a new regularization scheme accompanied with an efficient optimization procedure. On top of showing very competitive performance on artificial and real data sets, our method demonstrates capabilities for fine interpretation of its parameters, as illustrated in applications to genetics, genomics and spectroscopy.

1. Introduction. Multivariate regression, that is, regression with multiple response variables, is increasingly used to model high dimensional problems. By considering multiple responses, we wish to strengthen the estimation and/or the selection of the relevant input features, by taking advantage of the dependency pattern between the outputs. This is particularly appealing when the data is scarce, or even in the '$n < p$' high dimensional setup, in which framework this work enters. Compared to its univariate counterpart, the general linear model aims to predict several – say q – responses from a set of p predictors, relying on a training data set $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1,\dots,n}$:

$$\mathbf{y}_i = \mathbf{B}^T \mathbf{x}_i + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{\varepsilon}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \quad \forall i = 1, \dots, n. \tag{1}$$

The $p \times q$ matrix of regression coefficients $\mathbf{B}$ and the $q \times q$ covariance matrix $\mathbf{R}$ of the Gaussian noise $\boldsymbol{\varepsilon}_i$ are unknown. Model (1) has been studied by Mardia, Kent and Bibby (1979) in the low dimensional case where both ordinary and generalized least squares estimators of $\mathbf{B}$ coincide and do not

*This work was partially supported by the *GenSel* MetaProgram of the French National Institute for Research in Agronomy (INRA)

Keywords and phrases: Multivariate Regression, Regularization, Sparsity, Gaussian Graphical Model, Regulatory Motif, QTL study, Spectroscopy

depend on $\mathbf{R}$. These approaches boil down to perform q independent regressions, each column $\mathbf{B}_j$ describing the weights associating the p predictors to the jth response. In the $n < p$ setup however, these estimators are not defined.

The present work proposes a general multivariate regression framework suited to the high dimensional setup with three purposes that shall ease the learning task: *i*) We want to account for the dependency structure between the outputs, if it exists. That is to say, we want to integrate the estimation of $\mathbf{R}$ in the inference process. *ii*) In order to improve interpretability, we pay a prior attention to the direct links between predictors and responses. Indeed, according to the Gaussian graphical models (GGM), direct effects are measured by partial covariances between predictors and responses (denoted $\mathbf{\Omega_{xy}}$ hereafter), and not by the regression coefficients $\mathbf{B}$. In this regard, as few predictors are expected to *directly* influence each response, we choose to induce sparsity on partial covariances at once. *iii*) We want to have the possibility of integrating some prior information about the predictors. Indeed, in many applications such as genomics or spectral data analysis, the predictors display a structure (e.g. their position along the genome) and their influence on the responses are likely to display the same structure.

We design a new estimator to achieve the three aforementioned goals that we call SPRING (*"Structured selection of Primordial Relationships IN the General linear model"*). To illustrate the features of SPRING, we consider as a motivating example the cookie dough data from Osborne et al. (1984): in this typical multivariate regression problem, four continuous responses related to the composition of biscuits (percentages of fat, sucrose, flour and water) may be linearly explained by predictors corresponding to reflectance spectroscopy measurements. Parameters fitted by SPRING are represented on Figure 1 with, from left to right, the estimators of the regression coefficients $\mathbf{B}$, of the direct effects $\mathbf{\Omega_{xy}}$ and of the residual covariance $\mathbf{R}$. The regression coefficients show no sparsity pattern but a strong spatial structure along the spectrum, characterized by waves induced by a smooth first order difference prior. Concerning a potential structure between the outputs, we identify interesting regions where a strong correlation between the responses induces correlations between the regression coefficients. Consider for instance positions 1750 to 2000: the regression parameters related to 'dry-flour' are clearly anti-correlated with those related to 'sucrose'. Still, we cannot distinguish in $\hat{\mathbf{B}}$ for a direct effect of this region on either the flour or sucrose composition. Such a distinction is achieved on the middle panel where direct effects $\hat{\mathbf{\Omega}}_{\mathbf{xy}}$ selected by SPRING are plotted: it defines sparse predictive regions specific to each response, where it is now obvious that

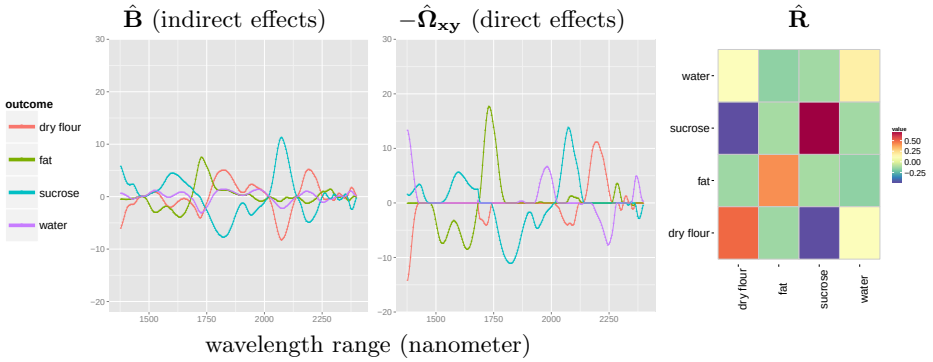

FIGURE 1. *Parameters estimated by SPRING for the cookies dough data (better seen in color). Since* $\hat{\mathbf{B}}$ *and* $\hat{\mathbf{\Omega}}_{\mathbf{xy}}$ *are opposite in sign in our model, we represent* $-\hat{\mathbf{\Omega}}_{\mathbf{xy}}$ *to ease the comparison between direct and indirect effects.*

region 1750 to 2000 is rather linked to the sucrose. This example continues latter, and also illustrates the good predictive performance of SPRING.

The outline of the paper is as follows. In Section 2 we expose the statistical model and the associated learning strategy: first, we introduce the multivariate regression model and the parametrization retained. Second, we propose a penalized likelihood criterion to regularized the model according to the three aforementioned goals. Then we develop an efficient optimization strategy to minimize this criterion. A paragraph also addresses the model selection issue. Section 3 is dedicated to illustrative simulation studies. In Section 4, we develop further the cookie dough example and investigate two multivariate problems: first, the relationships between genetic markers and a series of phenotypes of the plant *Brassica napus*; and second, the regulatory motifs discovery of yeast from time-course microarray experiments. In theses applications, specific underlying structuring priors arise, whose integration within our model is detailed.

Links to related works. Several recent papers tackle the multivariate regression problem with penalties-based approaches by including the dependency pattern of the outputs in the learning process. When this pattern is known *a priori*, we enter the "multitask" framework where many authors suggested variants of the ℓ_1-penalty shaped accordingly. A natural approach is to encourage a similar behavior of the regression parameters across the outputs – or tasks – using group-norms zeroing a full row of $\mathbf{B}$ at once, across the corresponding outputs (see e.g. Chiquet, Grandvalet and Ambroise, 2011; Obozinski, Wainwright and Jordan, 2011). Refinements exist to cope with complex dependency structures, using graph or tree structures (Kim and

Xing, 2009, 2010), yet the pattern remains fixed. When it is unknown, the general linear model (1) is a natural tool to account for dependencies between the responses. In this vein, Rothman, Levina and Zhu (2010); Yin and Li (2011) suggest to penalize the negative log-likelihood of (1) by two ℓ_1 norms respectively inducing sparsity on the regression coefficients $\mathbf{B}$ and on the inverse covariance $\mathbf{R}^{-1}$. However their criterion is hard to minimize as it is only bi-convex. Theoretical guarantees are proposed in Lee and Liu (2012), accompanied with a two-stage procedure involving a plug-in estimator of the inverse covariance, which leads to an optimization procedure computationally less demanding. In Sohn and Kim (2012), an elegant re-parametrization of (1) is depicted, leading to a formulation that is jointly convex and enjoys features of the GGM approach. Yuan and Zhang (2014) recently proposed theoretical analysis of this model. In this paper, we build on Sohn and Kim's parameterization yet we adopt a different stance regarding the regularization, both on the output structure and on the regression parameters. SPRING induces sparsity on the direct relationships like in Sohn and Kim (2012), yet no sparsity assumption is made for the inverse covariance, as we typically assume a reasonable number of outputs compared to the number of predictors and sample size. This leads to a different optimization strategy, a different estimator for the covariance $\mathbf{R}$ and eventually of the regression parameters $\mathbf{B}$. On top of that, we include structural information via an additional regularization term using an application-specific metric, generalizing the 'structured' versions of the Elastic-net (Slawski, zu Castell and Tutz, 2010; Hebiri and van De Geer, 2011) proposed in the univariate-output case. Indeed, an additional originality of our work is to investigate deeply the integration of such structural information in the various contexts of spectroscopy, genomic selection and regulatory motif discovery, by constructing our regularization scheme from each application specificity.

2. Model setup and learning.

2.1. *Convex parametrization of multivariate regression.* The statistical framework arises from a different parametrization of (1) which fills the gap between multivariate regression and GGM. Such a connection was first underlined in Sohn and Kim (2012), generalizing the link made by Witten and Tibshirani (2009) in the single output case between the LASSO and inverse covariance estimation. This amounts to investigate the joint probability distribution of $(\mathbf{x}_i, \mathbf{y}_i)$ in the Gaussian case, with the following block-wise decomposition of the covariance matrix $\mathbf{\Sigma}$ and its inverse $\mathbf{\Omega} = \mathbf{\Sigma}^{-1}$:

$$\mathbf{\Sigma} = \begin{pmatrix} \mathbf{\Sigma_{xx}} & \mathbf{\Sigma_{xy}} \\ \mathbf{\Sigma_{yx}} & \mathbf{\Sigma_{yy}} \end{pmatrix}, \quad \mathbf{\Omega} = \begin{pmatrix} \mathbf{\Omega_{xx}} & \mathbf{\Omega_{xy}} \\ \mathbf{\Omega_{yx}} & \mathbf{\Omega_{yy}} \end{pmatrix}. \tag{2}$$

Back to the distribution of $\mathbf{y}_i$ conditional on $\mathbf{x}_i$, multivariate Gaussian analysis shows that, for centered $\mathbf{x}_i$ and $\mathbf{y}_i$,

$$\mathbf{y}_i|\mathbf{x}_i \sim \mathcal{N}\left(-\mathbf{\Omega}_{\mathbf{yy}}^{-1}\mathbf{\Omega}_{\mathbf{yx}}\mathbf{x}_i, \mathbf{\Omega}_{\mathbf{yy}}^{-1}\right). \tag{3}$$

This model, associated to the full sample $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1,\dots,n}$, can be written in a matrix form by stacking in row the observations of the responses on the first hand, and of the predictors on the second hand, in two data matrices $\mathbf{Y}$ and $\mathbf{X}$ with respective sizes $n \times q$ and $n \times p$, so as

$$\mathbf{Y} = \mathbf{X}\mathbf{\Omega}_{\mathbf{xy}}\mathbf{\Omega}_{\mathbf{yy}}^{-1} + \boldsymbol{\varepsilon}, \qquad \mathrm{vec}(\boldsymbol{\varepsilon}) \sim \mathcal{N}\left(\mathbf{0}_{nq}, \mathbf{I}_p \otimes \mathbf{\Omega}_{\mathbf{yy}}^{-1}\right), \tag{4}$$

where, for a $n \times p$ matrix $\mathbf{A}$, denoting $\mathbf{A}_j$ its jth column, the vec operator is defined as $\mathrm{vec}(\mathbf{A}) = (\mathbf{A}_1^T \dots \mathbf{A}_p^T)^T$. The log-likelihood of (4) – which is a conditional likelihood regarding the joint model (2) – writes

$$\begin{aligned}\log L(\mathbf{\Omega}_{\mathbf{xy}}, \mathbf{\Omega}_{\mathbf{yy}}) = \frac{n}{2}\log|\mathbf{\Omega}_{\mathbf{yy}}| + \\ \frac{n}{2}\mathrm{tr}\left(\left(\mathbf{Y} - \mathbf{X}\mathbf{\Omega}_{\mathbf{xy}}\mathbf{\Omega}_{\mathbf{yy}}^{-1}\right)\mathbf{\Omega}_{\mathbf{yy}}\left(\mathbf{Y} - \mathbf{X}\mathbf{\Omega}_{\mathbf{xy}}\mathbf{\Omega}_{\mathbf{yy}}^{-1}\right)^T\right) + \mathrm{cst}.\end{aligned} \tag{5}$$

Introducing the empirical matrices of covariance $\mathbf{S}_{\mathbf{yy}} = n^{-1}\sum_{i=1}^n \mathbf{y}_i\mathbf{y}_i^T$, $\mathbf{S}_{\mathbf{xx}} = n^{-1}\sum_{i=1}^n \mathbf{x}_i\mathbf{x}_i^T$, and $\mathbf{S}_{\mathbf{yx}} = n^{-1}\sum_{i=1}^n \mathbf{y}_i\mathbf{x}_i^T$, one has

$$\begin{aligned}-\frac{2}{n}\log L(\mathbf{\Omega}_{\mathbf{xy}}, \mathbf{\Omega}_{\mathbf{yy}}) = -\log|\mathbf{\Omega}_{\mathbf{yy}}| + \mathrm{tr}\left(\mathbf{S}_{\mathbf{yy}}\mathbf{\Omega}_{\mathbf{yy}}\right) \\ + 2\mathrm{tr}\left(\mathbf{S}_{\mathbf{xy}}\mathbf{\Omega}_{\mathbf{yx}}\right) + \mathrm{tr}(\mathbf{\Omega}_{\mathbf{yx}}\mathbf{S}_{\mathbf{xx}}\mathbf{\Omega}_{\mathbf{xy}}\mathbf{\Omega}_{\mathbf{yy}}^{-1}) + \mathrm{cst}.\end{aligned} \tag{6}$$

This has been referred to as a 'conditional Gaussian Graphical Model' (cGGM) or 'partial Gaussian Graphical Model' in the recent literature (Yin and Li, 2011; Sohn and Kim, 2012; Lee and Liu, 2012; Yuan and Zhang, 2014). We notice by comparing the cGGM (3) to the multivariate regression model (1) that $\mathbf{\Omega}_{\mathbf{yy}}^{-1} = \mathbf{R}$ and $\mathbf{B} = -\mathbf{\Omega}_{\mathbf{xy}}\mathbf{\Omega}_{\mathbf{yy}}^{-1}$. Although equivalent to (1), the alternative parametrization (3) shows two important differences with several implications. First, in light of convex optimization theory, the negative log-likelihood (5) can be shown to be jointly convex in $(\mathbf{\Omega}_{\mathbf{xy}}, \mathbf{\Omega}_{\mathbf{yy}})$ (a formal proof can be found in Yuan and Zhang, 2014). Minimization problems involving (5) will thus be amenable to a global solution, which facilitates both optimization and theoretical analysis. As such, the conditional negative log-likelihood (5) will serve as a building block for our learning criterion. Second, it unveils new interpretations for the relationships between input and output variables, as discussed in Sohn and Kim (2012): $\mathbf{\Omega}_{\mathbf{xy}}$ describes the *direct* relationships between predictors and responses, the

support of which we are looking for to select relevant interactions. On the other hand, $\mathbf{B}$ entails *both direct and indirect* influences, possibly due to some strong correlations between the responses, described by the covariance matrix $\mathbf{R}$ (or equivalently its inverse $\mathbf{\Omega_{yy}}$). To provide additional insights on cGGM, Figure 2 illustrates the relationships between $\mathbf{B}, \mathbf{\Omega_{xy}}$ and $\mathbf{R}$ in two simple scenarios where $p = 40$ and $q = 5$. Scenarios a) and b) are discriminated by the presence of a strong structure among the predictors. Still, the important point to catch at this stage of the paper is how strong correlations between outcomes can completely "mask" the direct links in the regression coefficients: the stronger the correlation in $\mathbf{R}$, the less it is possible to distinguish in $\mathbf{B}$ the non-zero coefficients of $\mathbf{\Omega_{xy}}$.

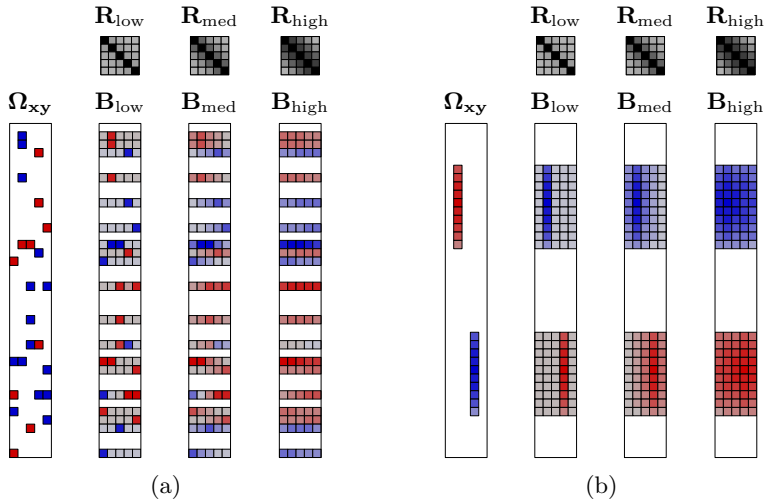

FIGURE 2. *Toy examples to illustrate the relationships between* $\mathbf{B}, \mathbf{\Omega_{xy}}$ *and* $\mathbf{R}$ *in the cGGM (better seen in colors): on panel a), a situation with no particular structure among the predictors; on panel b), a strong neighborhood structure. For each panel, we represent the effect of stronger correlations in* $\mathbf{R}$ *on masking the direct links in* $\mathbf{B}$.

2.2. *Structured regularization with underlying sparsity.* Our regularization scheme starts by considering some structural prior information as for the relationships between the coefficients. We are typically thinking of a situation where similar inputs are expected to have similar direct relationships with the outputs. Right panel of Figure 2 represents a cartoon of such a situation, where there exists an extreme neighborhood structure between the predictors. This depicts a pattern that acts along the rows of $\mathbf{B}$ or $\mathbf{\Omega_{xy}}$ as substantiated by the following Bayesian point of view.

Bayesian interpretation. Suppose the similarities can be encoded into a matrix $\mathbf{L}$. Our aim is to account for this information when learning the coefficients. The Bayesian framework provides a convenient setup to define the way the structural information should be accounted for. In the single output case (see, e.g. Marin and Robert, 2007), a conjugate prior for $\boldsymbol{\beta}$ would be $\mathcal{N}(\mathbf{0}, \mathbf{L}^{-1})$. Combined with the covariance between the outputs, this gives

$$\operatorname{vec}(\mathbf{B}) \sim \mathcal{N}(\mathbf{0}, \mathbf{R} \otimes \mathbf{L}^{-1}),$$

where $\otimes$ is the Kronecker product. By properties of the vec operator, this can be stated straightforwardly as

$$\operatorname{vec}(\mathbf{\Omega_{xy}}) \sim \mathcal{N}(\mathbf{0}, \mathbf{R}^{-1} \otimes \mathbf{L}^{-1})$$

for the direct links. Choosing such a prior results in

$$\log \mathbb{P}(\mathbf{\Omega_{xy}} | \mathbf{L}, \mathbf{R}) = \frac{1}{2} \operatorname{tr}\left(\mathbf{\Omega}_{\mathbf{xy}}^T \mathbf{L} \mathbf{\Omega_{xy}} \mathbf{R}\right) + \text{cst}.$$

Criterion. Through this argument, we propose a criterion with two terms to regularize the conditional negative log-likelihood (5): first, a smooth trace term relying on the available structural information $\mathbf{L}$; and second, a ℓ_1 norm that encourages sparsity among the direct links. We write the criterion as a function of $(\mathbf{\Omega_{xy}}, \mathbf{\Omega_{yy}})$ rather than $(\mathbf{\Omega_{xy}}, \mathbf{R})$, although equivalent in term of estimation, since the former leads to a convex formulation. The optimization problem turns to the minimization of

$$J(\mathbf{\Omega_{xy}}, \mathbf{\Omega_{yy}}) = - \frac{1}{n} \log L(\mathbf{\Omega_{xy}}, \mathbf{\Omega_{yy}}) + \frac{\lambda_2}{2} \operatorname{tr}\left(\mathbf{\Omega_{yx}} \mathbf{L} \mathbf{\Omega_{xy}} \mathbf{\Omega}_{\mathbf{yy}}^{-1}\right) + \lambda_1 \|\mathbf{\Omega_{xy}}\|_1. \tag{7}$$

Simple convexity arguments and requirements for positive definiteness of the Hessian of (7) lead to the following proposition, whose proof is omitted.

PROPOSITION 1. *The objective function* (7) *is jointly convex in* $(\mathbf{\Omega_{xy}}, \mathbf{\Omega_{yy}})$ *and admits at least one global minimum which is unique when* $n \geq q$ *and* $(\lambda_2 \mathbf{L} + \mathbf{S_{xx}})$ *is positive definite.*

2.3. *Connection to other sparse regression methods.* To get more insight about our model and to facilitate connections with existing approaches, we shall write the objective function (7) as a penalized *univariate* regression problem. This amounts to "vectorize" model (4) with respect to $\mathbf{\Omega_{xy}}$, i.e., to write the objective as a function of $(\boldsymbol{\omega}, \mathbf{\Omega_{yy}})$ where $\boldsymbol{\omega} = \operatorname{vec}(\mathbf{\Omega_{xy}})$. This

is stated in the following proposition, which can be derived from straightforward matrix algebra, as proved in Appendix Section A.1. The main interest of this proposition will become clear when deriving the optimization procedure to minimize (7), as the optimization problem when $\mathbf{\Omega_{yy}}$ is fixed turns to a generalized Elastic-Net problem. Remark that we use $\mathbf{A}^{1/2}$ to denote the square root of a matrix, obtained for instance by a Cholesky factorization in case of a symmetric positive definite matrix.

PROPOSITION 2. *Let* $\boldsymbol{\omega} = \mathrm{vec}(\mathbf{\Omega_{xy}})$. *An equivalent vector form of* (7) *is*

$$J(\boldsymbol{\omega}, \mathbf{\Omega_{yy}}) = -\frac{1}{2}\log|\mathbf{\Omega_{yy}}| + \frac{1}{2n}\left\|\tilde{\mathbf{y}} + \tilde{\mathbf{X}}\boldsymbol{\omega}\right\|_2^2 + \frac{\lambda_2}{2}\boldsymbol{\omega}^T\left(\mathbf{\Omega_{yy}^{-1}} \otimes \mathbf{L}\right)\boldsymbol{\omega} + \lambda_1\|\boldsymbol{\omega}\|_1,$$

where the $n \times q$ *dimensional vector* $\tilde{\mathbf{y}}$ *and the* $nq \times pq$ *dimensional matrix* $\tilde{\mathbf{X}}$ *depends on* $\mathbf{\Omega_{yy}}$ *such that*

$$\tilde{\mathbf{y}} = \mathrm{vec}(\mathbf{Y}\mathbf{\Omega_{yy}^{1/2}}), \quad \tilde{\mathbf{X}} = \left(\mathbf{\Omega_{yy}^{-1/2}} \otimes \mathbf{X}\right).$$

When there is only one response, that is, when $q = 1$, direct connections can be made with various existing approaches. In such cases, the variance-covariance matrix $\mathbf{R}$ and its inverse $\mathbf{\Omega_{yy}}$ turn to scalars, respectively denoted by σ^2 and ρ^2, with $\sigma = 1/\rho$. Applying Proposition 2, the objective (7) writes

$$J(\boldsymbol{\omega}, \rho) = -\log(\rho) + \frac{1}{2n}\|\rho\mathbf{y} + \mathbf{X}\boldsymbol{\omega}/(\rho)\|_2^2 + \frac{\lambda_2}{2\rho^2}\boldsymbol{\omega}^T\mathbf{L}\boldsymbol{\omega} + \lambda_1\|\boldsymbol{\omega}\|_1. \tag{8}$$

Its counterpart in term of the parametrization $(\boldsymbol{\beta}, \sigma)$, is

$$J(\boldsymbol{\beta}, \sigma) = \log(\sigma) + \frac{1}{2n\sigma^2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \frac{\lambda_2}{2\sigma^2}\boldsymbol{\beta}\mathbf{L}\boldsymbol{\beta} + \frac{\lambda_1}{\sigma^2}\|\boldsymbol{\beta}\|_1. \tag{9}$$

We recognize in (9) the log-likelihood of a "usual" linear model with Gaussian noise, unknown variance σ^2 and regression parameters $\boldsymbol{\beta}$, plus a mixture of ℓ_1 and ℓ_2 norm, a.k.a, an elastic-net penalization. This is the univariate analog of (1), whose criterion (8) is the convex reparametrizaton. Criterion (9) entail many regularization schemes found in the literature: assuming that σ is known, we recognize ridge regression when $\lambda_1 = 0$, the Lasso (Tibshirani, 1996) when $\lambda_2 = 0$, plus the Elastic-Net (Zou and Hastie, 2005) when $\mathbf{L} = \mathbf{I}_p$ and its extensions (Slawski, zu Castell and Tutz, 2010; Hebiri and van De Geer, 2011) which correspond to different choices of $\mathbf{L}$. Some natural choices for $\mathbf{L}$ motivated by applications in genetics, genomics and spectroscopy are discussed in Sections 3 and 4. If σ is unknown, letting $\lambda_2 = 0$ in (9) leads to the ℓ_1-penalized mixture regression model proposed

by Städler, Bühlmann and Geer (2010), where both $\boldsymbol{\beta}$ and σ^2 are inferred. In this work, Städler, Bühlmann and Geer insist on the importance of penalizing the vector of coefficients by an amount that is inversely proportional to σ. Penalizing the regression coefficients and the small variances σ^2 simultaneously is indeed highly desirable, which has some connections with the Bayesian Lasso of Park and Casella (2008).

To sum-up, our approach can be seen as a generalization of all these methods to the multiple-responses case, not to mentioned the structured regularization brought by the ℓ_2-norm.

2.4. *Optimization.* To minimize the objective, we rely on a simple strategy alternating the optimization on $\boldsymbol{\Omega}_{\mathbf{yy}}$ and $\boldsymbol{\Omega}_{\mathbf{xy}}$, that is,

$$\hat{\boldsymbol{\Omega}}_{\mathbf{yy}}^{(k+1)} = \arg\min_{\boldsymbol{\Omega}_{\mathbf{yy}} \succ 0} J_{\lambda_1\lambda_2}(\hat{\boldsymbol{\Omega}}_{\mathbf{xy}}^{(k)}, \boldsymbol{\Omega}_{\mathbf{yy}}), \tag{10a}$$

$$\hat{\boldsymbol{\Omega}}_{\mathbf{xy}}^{(k+1)} = \arg\min_{\boldsymbol{\Omega}_{\mathbf{xy}}} J_{\lambda_1\lambda_2}(\boldsymbol{\Omega}_{\mathbf{xy}}, \hat{\boldsymbol{\Omega}}_{\mathbf{yy}}^{(k+1)}). \tag{10b}$$

Alternate strategies were adopted by Rothman, Levina and Zhu (2010); Yin and Li (2011), yet in these cases the objective is only bi-convex as they rely on the classical parametrization (1) and their algorithms are not meant to converge to the global minimum. In Sohn and Kim (2012); Yuan and Zhang (2014), the criterion is jointly convex yet the regularization is different from ours and requires more demanding numerical optimization because of the regularization of the concentration matrix $\boldsymbol{\Omega}_{\mathbf{yy}}$. Besides, no convergence results are given in these works.

Concerning our criterion, the alternating scheme seems quite appropriate: first, we may apply the results developed in Tseng and Yun (2009); Tseng (2001) on the convergence of block coordinate descent for the minimization of nonsmooth separable function, since (7) is clearly separable in $(\boldsymbol{\Omega}_{\mathbf{xy}}, \boldsymbol{\Omega}_{\mathbf{yy}})$ for the nonsmooth part induced by the ℓ_1-norm. Thus, under the assumption of Proposition 1, the alternating scheme is guaranteed to convergence to the unique global minimum. Second, efficient procedures exist to solve the two convex sub-problems (10a) and (10b), as detailed herein.

On the first hand, (10b) can be recast as an Elastic-Net problem, which in turn, can be recast as a LASSO problem (see, e.g. Zou and Hastie, 2005; Slawski, zu Castell and Tutz, 2010). This is straightforward thanks to Proposition 2: when $\hat{\boldsymbol{\Omega}}_{\mathbf{yy}}$ is fixed, solution to (10b) can be obtained via

$$\hat{\boldsymbol{\omega}} \left(= \operatorname{vec}(\hat{\boldsymbol{\Omega}}_{\mathbf{xy}})\right) = \arg\min_{\boldsymbol{\omega} \in \mathbb{R}^{pq}} \frac{1}{2}\|\mathbf{A}\boldsymbol{\omega} - \mathbf{b}\|_2^2 + \frac{\lambda_2}{2}\boldsymbol{\omega}^T\tilde{\mathbf{L}}\boldsymbol{\omega} + \lambda_1\|\boldsymbol{\omega}\|_1, \tag{11}$$

where $\mathbf{A}, \mathbf{b}$ and $\tilde{\mathbf{L}}$ are defined by

$$\mathbf{A} = \left(\hat{\mathbf{\Omega}}_{\mathbf{yy}}^{-1/2} \otimes \mathbf{X}/\sqrt{n}\right), \mathbf{b} = -\text{vec}\left(\mathbf{Y}\hat{\mathbf{\Omega}}_{\mathbf{yy}}^{1/2}\right)/\sqrt{n} \text{ and } \tilde{\mathbf{L}} = \hat{\mathbf{\Omega}}_{\mathbf{yy}}^{-1} \otimes \mathbf{L}.$$

On the second hand, we can solve analytically (10a) with simple matrix algebra, as stated in the following Proposition whose proof is postponed to Appendix Section A.2.

PROPOSITION 3. *Let $n \geq q$. If $\hat{\mathbf{\Omega}}_{\mathbf{xy}} = \mathbf{0}$, $\hat{\mathbf{\Omega}}_{\mathbf{yy}} = \mathbf{S}_{\mathbf{yy}}^{-1}$. Otherwise, assume that the following eigen decomposition holds*

$$\hat{\mathbf{\Omega}}_{\mathbf{yx}}(\lambda_2\mathbf{L} + \mathbf{S}_{\mathbf{xx}})\hat{\mathbf{\Omega}}_{\mathbf{xy}}\mathbf{S}_{\mathbf{yy}} = \mathbf{U}\text{diag}(\boldsymbol{\zeta})\mathbf{U}^{-1}$$

and denote by $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_q)$ the respective unique positive roots of $\eta_j^2 - \eta_j - \zeta_j$. For a fix $\hat{\mathbf{\Omega}}_{\mathbf{xy}}$, solution $\hat{\mathbf{\Omega}}_{\mathbf{yy}}$ to (10a) and its inverse writes

$$\hat{\mathbf{\Omega}}_{\mathbf{yy}} = \mathbf{U}\text{diag}(\boldsymbol{\eta}/\boldsymbol{\zeta})\mathbf{U}^{-1}\hat{\mathbf{\Omega}}_{\mathbf{yx}}(\lambda_2\mathbf{L} + \mathbf{S}_{\mathbf{xx}})\hat{\mathbf{\Omega}}_{\mathbf{xy}}(= \hat{\mathbf{R}}^{-1}), \tag{12a}$$

$$\hat{\mathbf{\Omega}}_{\mathbf{yy}}^{-1} = \mathbf{S}_{\mathbf{yy}}\mathbf{U}\text{diag}(\boldsymbol{\eta}^{-1})\mathbf{U}^{-1}(= \hat{\mathbf{R}}). \tag{12b}$$

Because our procedure relies on alternating optimization, it is difficult to give either a global rate of convergence or a complexity bound. Nevertheless, the complexity of each iteration is easy to derive, since it amounts to two well known problems: the main computational cost in (10a) is due to the SVD of a $q \times q$ matrix, which costs $\mathcal{O}(q^3)$. Concerning (10b), it amounts to the resolution of an Elastic-Net problem with $p \times q$ variables and $n \times q$ samples. If the final number of nonzero entries in $\hat{\mathbf{\Omega}}_{\mathbf{xy}}$ is k, a good implementation with Cholesky update/downdate is roughly in $\mathcal{O}(npq^2k)$ (see, e.g. Bach et al., 2012). Since we typically assumed that $n \geq q$, the global cost of a single iteration of the alternating scheme is thus $\mathcal{O}(npq^2k)$, and we theoretically can treat problems with large p when k remains moderate.

Finally, we typically want to compute a series of solutions along the regularization path of Problem (7), that is, for various values of (λ_1, λ_2). To this end, we simply choose a grid of penalties $\Lambda_1 \times \Lambda_2 = \left\{\lambda_1^{\min}, \ldots, \lambda_1^{\max}\right\} \times \left\{\lambda_2^{\min}, \ldots, \lambda_2^{\max}\right\}$. The process is easily distributed on different computer cores, each core corresponding to a value picked in Λ_2. Then on each core – i.e. for a fix $\lambda_2 \in \Lambda_2$ – we cover all the values of $\lambda_1 \in \Lambda_1$ relying on the warm start strategy frequently used to go through the regularization path of ℓ_1-penalized problems. A prototype implementation is distributed in an `R`/`C++`-package available from the first author web site.

2.5. *Model selection and parameter tuning.* Model selection amounts here to choosing a couple of tuning parameters (λ_1, λ_2) among the grid $\Lambda_1 \times \Lambda_2$, where λ_1 tunes the sparsity while λ_2 tunes the amount of structural regularization. We aim either at picking the model with minimum prediction error, or the one closest to the true model assuming Equation (3) holds. These two perspectives generally do not lead to the same model choice: when looking for the model minimizing the prediction error, K-fold cross-validation is the recommended option (Hesterberg et al., 2008) despite its additional computational cost. Letting $\kappa : \{1, \ldots, n\} \to \{1, \ldots, K\}$ the function indexing the fold to which observation i is allocated, the CV-choice is

$$(\lambda_1^{\text{cv}}, \lambda_2^{\text{cv}}) = \underset{(\lambda_1,\lambda_2)\in\Lambda_1\times\Lambda_2}{\arg\min} \frac{1}{n}\sum_{i=1}^{n} \left\| \mathbf{x}_i^T \hat{\mathbf{B}}^{\lambda_1,\lambda_2}_{-\kappa(i)} - \mathbf{y}_i \right\|_2^2,$$

where $\hat{\mathbf{B}}^{\lambda_1,\lambda_2}_{-\kappa(i)}$ minimizes the objective (7), once data $\kappa(i)$ has been removed. In place of the prediction error, other quantities can be cross-validated in our context, e.g. the negative log-likelihood (5). In the remainder of the paper, we only consider cross-validating the prediction error, though.

As an alternative to cross-validation, analytic criteria provide a fast way to perform model selection and are sometimes more suited to the selection of the true underlying model. Although relying on asymptotic derivations, they often offer good practical performances when the sample size n is not too small compared to the problem dimension. For penalized methods, a general form for various information criteria is expressed as a function of the likelihood L (defined by (5) here) and the effective degrees of freedom:

$$(\lambda_1^{\text{pen}}, \lambda_2^{\text{pen}}) = \underset{(\lambda_1,\lambda_2)\in\Lambda_1\times\Lambda_2}{\arg\min} -2\log L(\hat{\mathbf{\Omega}}^{\lambda_1,\lambda_2}_{\mathbf{xy}}, \hat{\mathbf{\Omega}}^{\lambda_1,\lambda_2}_{\mathbf{yy}}) + \text{pen} \cdot \text{df}_{\lambda_1,\lambda_2}. \tag{13}$$

Setting pen $= 2$ or $\log(n)$ respectively lead to AIC or BIC. For the practical evaluation of (13), we must give some sense to the effective degrees of freedom $\text{df}_{\lambda_1,\lambda_2}$ of a SPRING fit. We use the now classical definition of Efron (2004) and rely on the work of Tibshirani and Taylor (2012) to derive the following Proposition, whose proof is postponed to Appendix Section A.3.

PROPOSITION 4. *An unbiased estimator of df_{λ_1,λ_2} for a SPRING fit is*

$$\hat{\text{df}}_{\lambda_1,\lambda_2} = \text{card}(\mathcal{A}) - \lambda_2 \text{tr}\left((\hat{\mathbf{R}} \otimes \mathbf{L})_{\mathcal{A}\mathcal{A}} \left((\hat{\mathbf{R}} \otimes (\mathbf{S}_{\mathbf{xx}} + \lambda_2 \mathbf{L}))_{\mathcal{A}\mathcal{A}} \right)^{-1} \right), \tag{14}$$

where $\mathcal{A} = \left\{ j : \text{vec}\left(\hat{\mathbf{\Omega}}^{\lambda_1,\lambda_2}_{\mathbf{xy}} \right) \neq 0 \right\}$ *is the set of nonzero entries in* $\hat{\mathbf{\Omega}}^{\lambda_1,\lambda_2}_{\mathbf{xy}}$.

Note that we can compute expression (14) at no additional cost, relying on computations already made during the optimization process.

3. Simulation studies. In this section, we would like to illustrate the new features of our proposal compared to several baselines in well controlled settings. To this end, we perform two simulation studies to evaluate i) the gain brought by the estimation of the residual covariance $\mathbf{R}$ and ii) the gain brought by the inclusion of informative structure on the predictors via $\mathbf{L}$.

Implementation details. In our experiments, performances are compared with well-established and distributed regularization methods: the LASSO (Tibshirani, 1996), the multitask group-LASSO, MRCE (Rothman, Levina and Zhu, 2010), the Elastic-Net (Zou and Hastie, 2005) and the Structured Elastic-Net (Slawski, zu Castell and Tutz, 2010). LASSO and group-LASSO are fitted with the `R`-package `glmnet` (Friedman, Hastie and Tibshirani, 2010) and MRCE with Rothman, Levina and Zhu's package. All other methods are fitted using our own code.

Data generation. Artificial datasets are generated according to the multivariate regression model (1). We assume that the decomposition $\mathbf{B} = \mathbf{\Omega_{xy}}\mathbf{\Omega_{yy}^{-1}} = \mathbf{\Omega_{xy}}\mathbf{R}$ holds for the regression coefficients. We control the sparsity pattern of $\mathbf{\Omega_{xy}}$ by arranging non null entries according to a possible structure of the predictors along the rows of $\mathbf{\Omega_{xy}}$. We always use uncorrelated Gaussian predictors $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ not to promote excessively the methods that take this structure into account. Strength of the relationships between the outputs are tuned by the covariance matrix $\mathbf{R}$. We measure the performance of the learning procedures thanks to the prediction error (PE) estimated using a large test set of observations generated according to the true model. When relevant, mean squared error (MSE) of the regression coefficients $\mathbf{B}$ is also presented. For conciseness, it is eluded when it shows quantitatively identical results as PE.

3.1. *Influence of covariance between outcomes.* The first simulation study aims to illustrate the advantage of splitting $\mathbf{B}$ into $-\mathbf{\Omega_{xy}}\mathbf{R}$ over working directly on $\mathbf{B}$. We set $p = 40$ predictors, $q = 5$ outcomes and randomly select 25 non null entries in $\{-1, 1\}$ to fill $\mathbf{\Omega_{xy}}$. We do not put any structure along the predictors as this study intends to measure the gain of using the cGGM approach. The covariance follows a Toeplitz scheme [1] so as $\mathbf{R}_{ij} = \tau^{|i-j|}$, for $i, j = 1, \ldots, q$. We consider three scenarios tuned by $\tau \in \{.1, .5, .9\}$ corresponding to an increasing correlation between the outcomes that eventually makes the cGGM more relevant. This settings have been used to generate panel (a) of Figure 2. For each covariance scenario, we generate a training

[1] We set $\mathbf{R}$ a correlation matrix not to excessively penalize the LASSO or the group-LASSO, that use the same tuning parameter λ_1 across the outcomes (and thus the same variance estimator).

set with size $n = 50$ and a test set with size $n = 1000$. We assess the performance of SPRING by comparison with three baselines: *i*) the LASSO, *ii*) the ℓ_1/ℓ_2 multitask group-LASSO and *iii*) SPRING with known covariance matrix $\mathbf{R}$. As it corresponds to the best fit that we can meet with our proposal, we call this variant the "oracle" mode of SPRING. The final estimators are obtained by 5-fold cross-validation on the training set. Figure 3 gives the boxplots of PE obtained for 100 replicates. As expected, the advantage of taking the covariance into account becomes more an more important to maintain a low PE when τ increases. When correlation is low, the LASSO dominates the group-LASSO; this is the other way round in the high correlation setup, where the latter takes full advantage of its grouping effect along the outcomes. Still, our proposal remains significantly better as soon as τ is substantial enough. We also note that our iterative algorithm does a good job since SPRING remains close to its oracle variant.

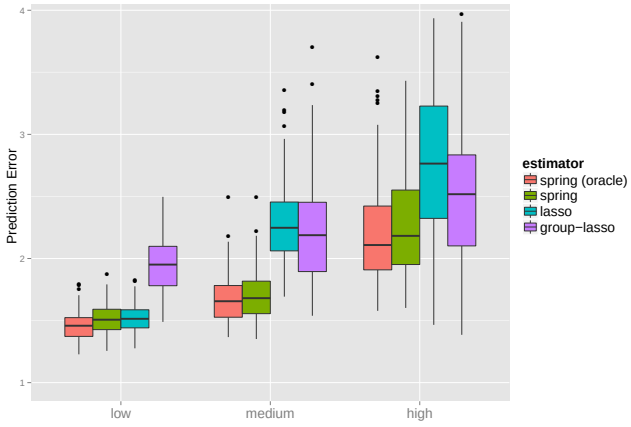

FIGURE 3. *Illustration of the influence of correlations between outcomes. Scenarios* $\{low, med, high\}$ *map to* $\tau \in \{.1, .5, .9\}$ *in the Toeplitz-shaped covariance matrix* $\mathbf{R}$.

3.2. *Structure integration and robustness.* The second simulation study is designed to measure the impact of introducing an informative prior along the predictors via $\mathbf{L}$. To remove any effect induced by the covariance between the outputs, we fix $q = 1$. In this case the criterion writes as in Expressions (8) and (9): $\mathbf{R}$ boils down to a scalar σ^2 and $\mathbf{B}, \mathbf{\Omega_{xy}}$ turn to two p-size vectors sharing the same sparsity pattern so as $\boldsymbol{\beta} = -\boldsymbol{\omega}\sigma^2$. In this situation, SPRING is close to Slawski, zu Castell and Tutz's structured Elastic-Net, except that we hope for a better estimation of the coefficients thanks to the estimation of σ. For comparison, we thus draw inspiration from the simulation settings originally used to illustrate the structured Elastic-Net: we set $\boldsymbol{\omega} = (\omega_j)_{j=1}^p$, with $p = 100$, so as we observe a sparse vector with two

smooth bumps, one positive and the other one negative:

$$\omega_j = \begin{cases} -((30-j)^2 - 100)/200 & j = 21, \dots 39, \\ \quad ((70-j)^2 - 100)/200 & j = 61, \dots 80, \\ 0 & \text{otherwise.} \end{cases}$$

A natural choice for $\mathbf{L}$ is to rely on the first forward difference operator as in the fused-Lasso (Tibshirani et al., 2005), or it smooth counterpart (Hebiri and van De Geer, 2011). We thus set $\mathbf{L} = \mathbf{D}^T\mathbf{D}$ with

$$\mathbf{D}_{ij} = \begin{cases} 1 & \text{if } i = j, \\ -1 & \text{if } j = i+1\ , \\ 0 & \text{otherwise} \end{cases} \qquad \begin{array}{c} i = 1, \dots, p-1, \\ j = 1, \dots, p. \end{array} \tag{15}$$

Hence, $\mathbf{L}$ is the combinatorial Laplacian of a chain graph between the successive predictors. Figure 4 shows the typical gain brought by prior structure

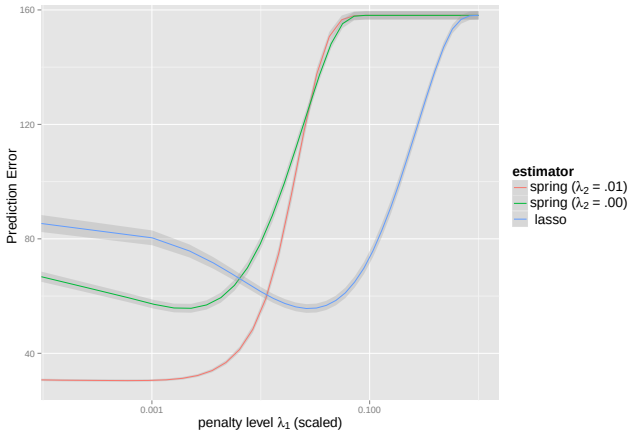

FIGURE 4. *Illustrating gain brought by structure information toward predictive performance*

knowledge. We generate 120 samples for the learning with $\sigma^2 = 5$ and represent 100 averaged PE curves as a function of λ_1 for the LASSO and for two versions of SPRING, with ($\lambda_2 = .01$) and without ($\lambda_2 = 0$) informative prior. Incorporating a relevant structural information leads to a dramatic improvement. As expected, univariate SPRING with no prior performs like the LASSO, in the sense that they share the same minimal PE.

Now, the question is: what if we introduce a wrong prior, that is, a matrix $\mathbf{L}$ completely disconnected from the true structure of the coefficients? To answer this question, we use the same settings as above and randomly swap

the entries in $\boldsymbol{\omega}$, using exactly the same $\mathbf{x}_i$ to generate the $\mathbf{y}_i$[2]. We then apply SPRING using respectively a non informative prior equal to the identity matrix (thus mimicking the Elastic-Net) and a 'wrong' prior $\mathbf{L}$ whose rows and columns remain unswapped. We also try the LASSO, the Elastic-Net and the structured Elastic-Net. All methods are tuned with 5-fold cross-validation on the learning set. Table 1 presents the results averaged over 100 runs both in term of PE (using a size-1000 test set) and MSE.

Method	*Scenario*	*MSE*	*PE*
LASSO	–	.336 (.096)	58.6 (10.2)
E-Net ($\mathbf{L} = \mathbf{I}$)	–	.340 (.095)	59 (10.3)
SPRING ($\mathbf{L} = \mathbf{I}$)	–	.358 (.094)	60.7 (10)
S. E-net	unswapped	.163 (.036)	41.3 (4.08)
($\mathbf{L} = \mathbf{D}^T\mathbf{D}$)	swapped	.352 (.107)	60.3 (11.42)
SPRING	unswapped	.062 (.022)	31.4 (2.99)
($\mathbf{L} = \mathbf{D}^T\mathbf{D}$)	swapped	.378 (.123)	62.9 (13.15)

TABLE 1

Structure integration: performance and robustness. Swapped and unswapped results are the same for LASSO, Elastic-Net and SPRING for $\mathbf{L} = \mathbf{I}$.

As expected, the methods that do not integrate any structural information (LASSO, Elastic-Net and SPRING with $\mathbf{L} = \mathbf{I}$) are not affected by the permutation, and we avoid these redundancies in Table 1 to save space. Overall, they share similar performance both in term of PE and MSE. When the prior structure is relevant, SPRING, and to a lesser extent the structured Elastic-Net, clearly outperform the other competitors. Surprisingly, this is particularly true in term of MSE, where SPRING also dominates the structured Elastic-Net that disposes of the same information. This means that the estimation of the variance also helped in the inference process. Finally, these results essentially support the robustness of the structured methods which are not much altered when using a wrong prior specification.

4. Application Studies. This Section aims at assessing predictive performance and eventually interpretability of our approach for real application purposes. We believe that such an aim is hard to achieve on artificial data, where the choice of neutral settings is difficult. This is particularly true for our model which involves integration of structural information that is deeply related to the problem at hand. Thus, to demonstrate the flexibility of SPRING, we investigate three multivariate data problems from various contexts, namely spectroscopy, genetics and genomics, where we insist on the construction of the structuring matrix $\mathbf{L}$.

[2]We also used the same seed and CV-folds for both scenarios.

4.1. *Near-Infrared Spectroscopy of Cookie Dough Pieces.*

Context. In Near-Infrared (NIR) spectroscopy, one aims to predict one or several quantitative variables from the NIR spectrum of a given sample. Each sampled spectrum is a curve that represents the level of reflectance along the NIR region, that is, wavelengths from 800 to 2500 nanometers (nm). The quantitative variables are typically related to the chemical composition of the sample. The problem is then to select the most predictive region of the spectrum, that is, some peaks that show good capabilities for predicting the response variable(s). This is known as a "calibration problem" in Statistics. NIR technique is used in fields as diverse as agronomy, astronomy or pharmacology. In such experiments, it is likely to meet very strong correlations and structure along the predictors. In this perspective, Hans (2011) proposes to apply the Elastic-Net which is known to selected simultaneously groups of correlated predictors. However it is not adapted to the prediction of several responses simultaneously. In Brown, Fearn and Vannucci (2001), an interesting wavelets regression model with Bayesian inference is introduced that enters the multivariate regression model, just like our proposal.

Description of the dataset. Continuing the example mentioned in the introduction, we consider the cookie dough data from Osborne et al. (1984). The data with the corresponding test and training sets are available in the **fds** `R` package. After data pretreatments as in Brown, Fearn and Vannucci (2001), we have $n = 39$ dough pieces in the training set: each sample consists in an NIR spectrum with $p = 256$ points measured from 1380 to 2400 nm (spaced by 4 nm), and in four quantitative variables that describe the percentages in fat, sugar flour and water of the dough piece.

Structure specification. We would like to account for the neighborhood structure between the predictors which is obvious in the context of NIR spectroscopy: since spectrum are continuous curves, a smooth neighborhood prior will encourage predictors to be selected by "wave", which seems more satisfactory than isolated peaks. Thus, we naturally define $\mathbf{L}$ by means of the first forward difference operator (15). We also tested higher orders of the operator to induce a stronger smoothing effect, but they do not lead to dramatic changes in term of PE and we omit them. Order k is simply obtained by powering the first order matrix. Such techniques have been studied in a structured ℓ_1-penalty framework in Kim et al. (2009); Tibshirani and Taylor (2011) and is known as *trend filtering*. Our approach is different, though: it enters a multivariate framework and it is based on a smooth ℓ_2-penalty coupled with the ℓ_1-penalty for selection.

Results. The predictive performances of a series of regression techniques are compared in Table 2, evaluated on the same test set. The first four

Method	*fat*	*sucrose*	*flour*	*water*
Stepwise MLR	.044	1.188	.722	.221
Decision theory	.076	.566	.265	.176
PLS	.151	.583	.375	.105
PCR	.160	.614	.388	.106
Wavelet Regression	.058	.819	.457	.080
LASSO	.044	.853	.370	.088
group-LASSO	.127	.918	.467	.102
MRCE	.151	.821	.321	.081
Structured E-net	.044	.596	.363	.082
SPRING (CV)	.065	.397	.237	.083
SPRING (BIC)	.048	.389	.243	.066

TABLE 2

Prediction error for the cookie dough data.

rows (namely stepwise multivariate regression, Bayesian decision theory approach, Partial Least Square and Principal Component Regression) correspond to calibration methods originally presented in Brown, Fearn and Vannucci (2001); row five (Bayesian Wavelet Regression) is Brown, Fearn and Vannucci's original proposal; the remainder of the table is due to our own analysis. We observe that, on this example, SPRING achieves extremely competitive results. The BIC performs notably well as a model selection criterion. Figure 5 shows the regression coefficients adjusted with the penal-

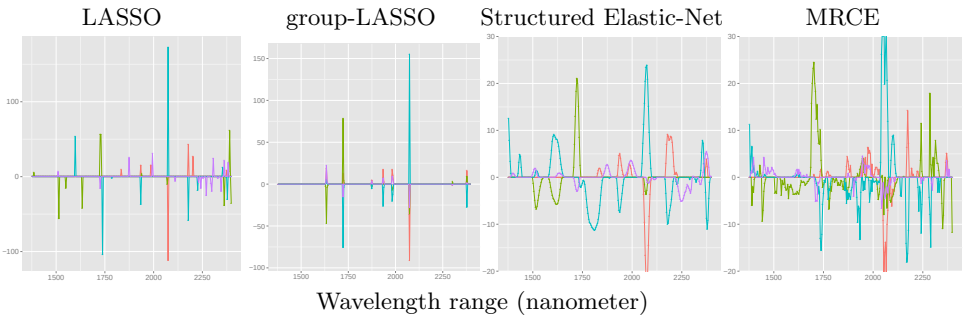

FIGURE 5. *Parameters estimated by the penalized regression methods for the cookie data*

ized regression techniques: apart from LASSO which selects very isolated (and unstable) wavelengths, non-zero regression coefficients are very spread, and therefore hard to interpret. As expected, the multitask group-Lasso activates the same predictors across the responses, which is not a good idea

when looking at the predictive performance in Table 2. On the other hand, back to the beginning of the paper in Figure 1, the direct effects selected by SPRING with BIC define predictive regions specific of each response which are well suited for interpretability purpose.

4.2. *Multi-trait Genomic Selection in Brassica napus.*

Context. Genomic selection aims at predicting one or several phenotypes based on the information of genetics markers. In plant and animal genetics, deriving an accurate prediction of complex traits is of major importance for the early detection and selection of individuals of high genetic value. Due to the continued advancement of high-throughput genotyping and sequencing technologies, genetic information is now available at low cost, whereas acquiring trait information remains expensive. A typical experiment will only contain a few hundreds of phenotyped individuals genotyped for thousands of markers. Consequently regularization methods such as ridge or Lasso regression or their Bayesian counterparts have been proposed since the very beginning of genomic selection (de los Campos et al., 2012). Still, in most studies only single trait genomic selection is performed, neglecting correlations between phenotypes. Moreover, little attention has been devoted to the development of regularization methods including prior genetic knowledge.

Description of the dataset. We consider the *Brassica napus* dataset described in Ferreira et al. (1995) and Kole et al. (2002). Data consists in $n = 103$ double-haploid lines derived from 2 parent cultivars, 'Stellar' and 'Major', on which $p = 300$ genetic markers and $q = 8$ traits (responses) were recorded. Each marker is a 0/1 covariate with $x_i^j = 0$ if line i has the 'Stellar' allele at marker j, and $x_i^j = 1$ otherwise. Traits included are percent winter survival for 1992, 1993, 1994, 1997 and 1999 (surv92, surv93, surv94, surv97, surv99, respectively), and days to flowering after no vernalization (flower0), 4 weeks vernalization (flower4) or 8 weeks vernalization (flower8).

Structure specification. In a biparental line population, correlation between 2 markers depends on their genetic distance defined in terms of recombination fraction. As a consequence, one expects adjacent markers on the sequence to be correlated, yielding similar direct relationships with the phenotypic traits. Noting d_{12} the genetic distance between markers M_1 and M_2, one has $\text{cor}(M_1 1, M_2 2) = \rho^{d_{12}}$, where $\rho = .98$[3]. The covariance matrix $\mathbf{L}^{-1}$ can hence be defined as $\mathbf{L}^{-1}_{ij} = \rho^{d_{ij}}$. Moreover, assuming recombination events are independent between M_1 and M_2 on one hand, and M_2 and M_3

[3]This value directly arises from the definition of the genetic distance itself.

on the other hand, one has $d_{13} = d_{12} + d_{23}$ and matrix $\mathbf{L}^{-1}$ exhibits an inhomogeneous AR(1) profile. As a consequence, $\mathbf{L}$ is tridiagonal with general elements

$$\begin{aligned} w_{i,i} &= \frac{1-\rho^{2d_{i-1,i}+2d_{i,i+1}}}{(1-\rho^{2d_{i-1,i}})(1-\rho^{2d_{i,i+1}})}, \\ w_{i,i+1} &= \frac{-\rho^{d_{i,i+1}}}{1-\rho^{2d_{i,i+1}}} \end{aligned}$$

and $w_{i,j} = 0$ if $|i-j| > 1$. For the first (resp. last) marker, the distance $d_{i-1,i}$ (resp. $d_{i,i+1}$) is infinite.

Results. We apply SPRING to study the influence of each marker on the traits and compare its predictive performance with these of its competitors. PE is estimated by randomly splitting the 103 samples into training and test sets with sizes 93 and 10. Before adjusting the models, we first scale the outcomes on the training and test sets to facilitate interpretability. Five-fold cross-validation is used on the training set to choose the tuning parameters. Two hundred random samplings of the test and training sets were conducted to estimate the PE given in Table 3. All methods provide similar results

Method	*surv92*	*surv93*	*surv94*	*surv97*	*surv99*
LASSO	.730 (.011)	.977 (.009)	.943 (.010)	.947 (.009)	.916 (.010)
S. Enet	**.697** (.011)	.987 (.009)	.941 (.011)	.945 (.009)	.911 (.010)
MRCE	.759 (.010)	**.919** (.003)	917 (.006)	**.924** (.004)	.926 (.006)
SPRING	.724 (.010)	.948 (.008)	**.848** (.010)	.940 (.006)	**.907** (.009)

Method	*flower0*	*flower4*	*flower8*
LASSO	.609 (.011)	.501 (.011)	.744 (.011)
S. Enet	.577 (.011)	.478 (.010)	.727 (.012)
MRCE	.591 (.011)	.479 (.011)	.736 (.011)
SPRING	**.489** (.010)	**.419** (.009)	**.616** (.012)

TABLE 3

Estimated prediction error for the Brassica napus *data (standard error in parenthesis)*

although SPRING provides the smallest error for half of the traits. A picture of the between-response covariance matrix estimated with SPRING is given in Figure 6. In this context, this matrix reflects the correlation between the traits that are either explained by an unexplored part of the genotype, by the environment or by some interaction between the two. The residuals of the flowering times exhibit strong correlations, whereas the correlation between the survival rates are weak. It also shows that the survival traits have a larger residual variability than the flowering traits, suggesting a higher sensitivity to environmental conditions.

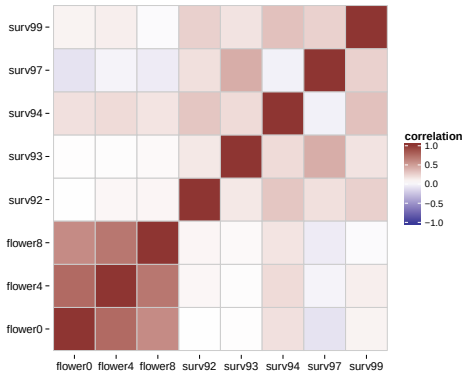

FIGURE 6. *Brassica study: residual covariance estimation*

We then turn to the effects of each marker on the different traits. The left panels of Figure 7 gives both the regression coefficients (top) and the direct effects (bottom). The gray zones correspond to chromosomes 2, 8 and 10, respectively. The exact location of the markers within these chromosomes are displayed in the right panels, where the size of the dots reflects the absolute value of the regression coefficients (top) and of the direct effects (bottom). The interest of considering direct effects rather than regression coefficients

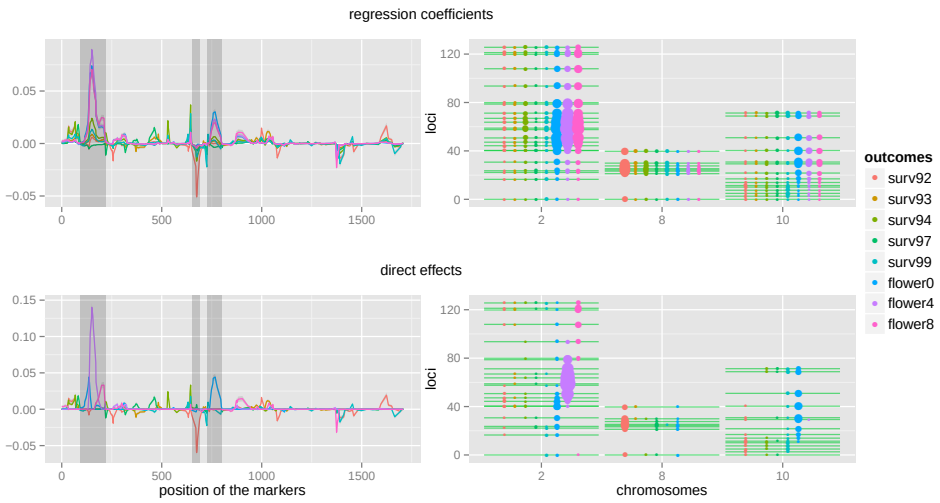

FIGURE 7. *Brassica Study: direct and indirect genetic effects of the markers on the traits estimated by SPRING (better seen in color).*

appears clearly here, looking for example at chromosome 2. Three large

overlapping regions are observed in the coefficient plot, for each flowering trait. A straightforward interpretation would suggest that the corresponding region controls the general flowering process. The direct effect plot allows to go deeper and shows that these three responses are actually controlled by three separated sub-regions within this chromosome. The confusion in the coefficient plot only results from the strong correlations observed between the three flowering traits.

4.3. *Selecting regulatory motif from multiple microarrays.*

Context. In genomics, the expression of genes is initiated by transcription factors that bind to the DNA upstream from the coding regions, called *regulatory regions*. This binding occurs when a given factor recognizes a certain (small) sequence called transcription factor binding site or *regulatory motif*. The presence of a regulatory motif in the regulatory region of a gene is therefore required to launch the transcriptional machinery. Conversely, genes hosting the same regulatory motif will be jointly expressed under certain conditions. As the binding relies on chemical affinity, some degeneracy can be tolerated in the motif definition and motifs similar up to small variations may share the same functional properties (see, e.g. Lajoie et al., 2012).

We are interested in the detection of such regulatory motifs, the presence of which controls genes expression profile. To this aim we try to establish a relation between the expression level of all genes across a series of conditions with the content of their respective regulatory regions in terms of motifs. In this context, we expect *i*) the set of influential motifs to be small for each condition, *ii*) the influent motifs for a given condition to be degenerated versions of each other, and *iii*) the expression under similar conditions to be controlled by the same motifs.

Description of the dataset. In Gasch et al. (2000), a series of microarray experiments are conducted on yeast cells (*Saccharomyces cerevisae*). Among these assays, we consider 12 time-course experiments profiling $n = 5883$ genes under various environmental changes as listed in Table 4. These expression sets form 12 potential response matrices $\mathbf{Y}$ whose number of column corresponds to the number of time points. Concerning the predictors, we consider the set of all motifs with length k formed with the four nucleotides, that is $\mathcal{M}_k = \{A, C, G, T\}^k$. There are $p = |\mathcal{M}_k| = 4^k$ such motifs. Unless otherwise stated, the motifs in $\mathcal{M}$ are lined up in lexicographical order e.g., when $k = 2$, $AA, AC, AG, AT, CA, CC, \ldots$ and so on. Then, the $n \times p$ matrix of predictors $\mathbf{X}$ is filled such as X_{ij} equals the occurrence count of motif j in the regulatory region of gene i.

Experiment	*# time point*	*# motifs selected by SPRING*		
		$k = 7$	$k = 8$	$k = 9$
Heat shock	8	30	68	43
Shift from 37°to 25 °C	5	3	11	33
Mild Heat shock	4	24	13	23
Response to H_2O_2	10	15	10	21
Menadione Exposure	9	16	1	7
DDT exposure 1	8	15	10	30
DDT exposure 2	7	11	33	21
Diamide treatment	8	45	25	35
Hyper-osmotic shock	7	36	24	15
Hypo-osmotic shock	5	20	8	29
Amino-acid starvation	5	47	30	39
Diauxic Shift	7	16	14	20
total number of unique motif inferred		87	82	72

TABLE 4

Time-course data from Gasch et al. (2000) considered for regulatory motif discovery

Structure specification. As we expect influential motifs for a given condition to be degenerated versions of each other, we would like to take this information into account in the matrix $\mathbf{L}$. To do so, we first measure the similarity between any two motifs from $\mathcal{M}_k$ with the Hamming defined as

$$\forall a, b \in \mathcal{M}_k, \quad \text{dist}(a, b) = \text{card}\{i : a_i \neq b_i\}.$$

For a distance of interest $0 \leq \ell \leq k$, we further define the $m \times m$ *symmetric* matrix $\mathbf{D}^{k,\ell} = (d_{ab}^{k,\ell})_{a,b \in \mathcal{M}_k}$ so as each entry $d_{ab}^{k,\ell}$ indicates if the Hamming distance is less or equal to ℓ. We refer to this distance for fixed value of ℓ as the *ℓ-distance matrix*:

$$\forall a, b \in \mathcal{M}_k, \quad d_{ab}^{k,\ell} = \begin{cases} 1 & \text{if dist}(a, b) \leq \ell \\ 0 & \text{otherwise.} \end{cases}$$

For motifs in $\mathcal{M}_k$ lined up in lexicographical order, manipulation of $\mathbf{D}$ greatly simplifies as motifs can be mapped to the set of natural integers $\mathbb{N}$. The following proposition shows that $\mathbf{D}^{k,\ell}$ can be computed recursively for increasing values of k, ℓ at the cost of a simple Kronecker product.

PROPOSITION 5. *For $\ell \leq k$, and considering that*

$$\mathbf{D}^{1,0} = \begin{array}{c} \\ A \\ C \\ G \\ T \end{array} \begin{array}{c} \begin{array}{cccc} A & C & G & T \end{array} \\ \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix} \end{array}, \quad \mathbf{D}^{1,1} = \begin{array}{c} \\ A \\ C \\ G \\ T \end{array} \begin{array}{c} \begin{array}{cccc} A & C & G & T \end{array} \\ \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \end{array},$$

then the ℓ-distance matrix between motifs of size $k+1$ can be expressed as a function of distance between motifs of size k as

$$\mathbf{D}^{k+1,\ell} = \left(\mathbf{D}^{1,0} \otimes \mathbf{D}^{k,\ell}\right) \vee \left(\mathbf{D}^{1,1} \otimes \mathbf{D}^{k,\ell-1}\right),$$

where $\vee$ is the OR operator and $\ell \geq 1$. Similarly, for values of $\ell : \ell + 1 \leq k$,

$$\mathbf{D}^{k,\ell+1} = \left(\mathbf{D}^{1,0} \otimes \mathbf{D}^{k-1,\ell}\right) \vee \left(\mathbf{D}^{1,1} \otimes \mathbf{D}^{k,\ell}\right).$$

$\mathbf{D}^{k,\ell}$ can be viewed as the adjacency matrix of a graph where the nodes are the motifs and where an edge is present between 2 nodes when the 2 motifs are at a Hamming distance less or equal to ℓ. We finally use the Laplacian of this graph as a structuring matrix $\mathbf{L}^{k,\ell} = (\ell_{ab}^{k,\ell})_{a,b\in\mathcal{M}_k}$, that is

$$\ell_{ab}^{k,\ell} = \begin{cases} \sum_{c\in\mathcal{M}_k} d_{ac}^{k,\ell} & \text{if } a = b, \\ -1 & \text{if } d_{ab}^{k,\ell} = 1, \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

Note that we came across a similar proposal by Li et al. (2010) in the context of sparse, structured PLS. The derivation is different though, and the objective of this method is not the selection of motifs but of family of motifs via the compression performed by the PLS.

Results. We apply our methodology for candidate motifs from $\mathcal{M}_7, \mathcal{M}_8$ and $\mathcal{M}_9$, which result in three lists of putative motifs having a direct effect on gene expression. Due to the very large number of potential predictors that comes with a sparser matrix $\mathbf{X}$ when k increases, we first perform a screening step that keep the $5,000$ motifs with the highest marginal correlations with $\mathbf{Y}$. Second, SPRING is fitted on each of the twelve time-course experiments described in Table 4. The selection of (λ_1, λ_2) is performed on a grid using the BIC (13) rather than the cross-validation approach, since we clearly face a variable selection problem. At the end of the day, the three lists corresponding to $k = 7, 8, 9$ comprehends respectively 87, 82 and 72 motifs, for which at least one coefficient in the associated row $\widehat{\mathbf{\Omega}}_{\mathbf{xy}}(j, \cdot)$ was found non-null for some of the twelve experiments, as detailed in Table 4.

To assess the relevance of the selected motifs, we compared them with the `MotifDB` patterns available in Bioconductor (Shannon (2013)), where known transcription factor binding sites are recorded. There are 453 such reference motifs with size varying from 5 to 23 nucleotides. Consider the case of $k = 7$ for instance: among the 87 SPRING motifs, 62 match one `MotifDB` pattern each and 25 are clustered into 11 `MotifDB` patterns as depicted in Table 5.

CTAAGCCAC TAGCCCC GCGCCCC	GCATGTGAA CCATATG TTGTGAG	CATGTAATT TGTAAAT TGTATAT	TGAAACA TTAGACC TAAAAAG
TGATCGGCGCCGCACGACGA GTATAAC GCGCCGT	TGCTGGTT GCTGGTT GCTGGTG	GATCGTATGATA ATCATAT TTGGTAT	ACGCGAAAA AACGAAA ACGAAAA
CCATACATCAC CATAGAC ATATCAC	ATTGACCTGGTC TCGACTT CGACTTG CCAGCTT	GACTAGATATATATATTCGAT ATATATT CATATAT ATATATG ATATATA	

TABLE 5

Comparison of SPRING selected motifs with `MotifDB` *patterns. Each cell corresponds to a* `MotifDB` *pattern (top) compared to a set of aligned SPRING motifs with size 7 (down).*

As seen in this table, the clusters of motifs selected by SPRING correspond to sets of variants of a same pattern. These clusters consist of motifs that are close according to the similarity encoded in the structure matrix $\mathbf{L}$. In this example, the framework we propose enabled us to orient the regression problem to account for motif degeneracy and helped us to select motifs that are consistent known binding sites. This example illustrates the flexibility of SPRING, which allows to use domain-specific definitions of the structure between the predictors.

Acknowledgments. We would like to thank Mathieu Lajoie and Laurent Bréhélin for kindly sharing the Gasch et al.'s dataset.

APPENDIX A: PROOFS

Most of the proofs rely on basic algebra and properties of the trace and the vec operators (see e.g. Harville, 1997), in particular

$$\begin{aligned}\operatorname{tr}(\mathbf{A}^T\mathbf{B}\mathbf{C}\mathbf{D}^T) &= \operatorname{vec}(\mathbf{A})^T(\mathbf{D}\otimes\mathbf{B})\operatorname{vec}(\mathbf{C}),\\ \operatorname{vec}(\mathbf{A}\mathbf{B}\mathbf{C}) &= (\mathbf{C}^T\otimes\mathbf{A})\operatorname{vec}(\mathbf{B}),\\ (\mathbf{A}\otimes\mathbf{B})(\mathbf{C}\otimes\mathbf{D}) &= (\mathbf{A}\mathbf{C})\otimes(\mathbf{C}\mathbf{D}).\end{aligned}$$

A.1. Derivation of Proposition 2. Concerning the two regularization terms in (7), we have $\|\boldsymbol{\Omega}_{\mathbf{xy}}\|_1 = \|\boldsymbol{\omega}\|_1$ since the ℓ_1-norm applies element-wise here and

$$\operatorname{tr}(\boldsymbol{\Omega}_{\mathbf{yx}}\mathbf{L}\boldsymbol{\Omega}_{\mathbf{xy}}\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}) = \boldsymbol{\omega}^T(\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}\otimes\mathbf{L})\boldsymbol{\omega}.$$

As for the log-likelihood (5), we work on the trace term:

$$\begin{aligned}
&\operatorname{tr}\left(\left(\mathbf{Y}+\mathbf{X}\boldsymbol{\Omega}_{\mathbf{xy}}\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}\right)^{\mathsf{T}}\left(\mathbf{Y}+\mathbf{X}\boldsymbol{\Omega}_{\mathbf{xy}}\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}\right)\boldsymbol{\Omega}_{\mathbf{yy}}\right)\\
&\quad= \operatorname{vec}\left(\mathbf{Y}+\mathbf{X}\boldsymbol{\Omega}_{\mathbf{xy}}\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}\right)^{T}\left(\boldsymbol{\Omega}_{\mathbf{yy}}\otimes\mathbf{I}_n\right)\operatorname{vec}\left(\mathbf{Y}+\mathbf{X}\boldsymbol{\Omega}_{\mathbf{xy}}\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}\right)\\
&\quad= \left\|\left(\boldsymbol{\Omega}_{\mathbf{yy}}\otimes\mathbf{I}_n\right)^{1/2}\left(\operatorname{vec}(\mathbf{Y})+\left(\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}\otimes\mathbf{X}\right)\operatorname{vec}(\boldsymbol{\Omega}_{\mathbf{xy}})\right)\right\|_2^2\\
&\quad= \left\|\operatorname{vec}(\mathbf{Y}\boldsymbol{\Omega}_{\mathbf{yy}}^{1/2})+\left(\boldsymbol{\Omega}_{\mathbf{yy}}^{-1/2}\otimes\mathbf{X}\right)\boldsymbol{\omega}\right\|_2^2.
\end{aligned}$$

The rest of the proof is straightforward.

A.2. Derivation of Proposition 3. By differentiation of the objective (7) over $\boldsymbol{\Omega}_{\mathbf{yy}}^{-1}$ we obtain the quadratic form

$$\boldsymbol{\Omega}_{\mathbf{yy}}\mathbf{S}_{\mathbf{yy}}\boldsymbol{\Omega}_{\mathbf{yy}}-\boldsymbol{\Omega}_{\mathbf{yy}}=\boldsymbol{\Omega}_{\mathbf{yx}}(\lambda_2\mathbf{L}+\mathbf{S}_{\mathbf{xx}})\boldsymbol{\Omega}_{\mathbf{xy}}.$$

After right multiplying both sides by $\mathbf{S}_{\mathbf{yy}}$, it becomes obvious that $\boldsymbol{\Omega}_{\mathbf{yy}}\mathbf{S}_{\mathbf{yy}}$ and $\boldsymbol{\Omega}_{\mathbf{yx}}(\lambda_2\mathbf{L}+\mathbf{S}_{\mathbf{xx}})\boldsymbol{\Omega}_{\mathbf{xy}}\mathbf{S}_{\mathbf{yy}}$ commute and thus share the same eigenvectors $\mathbf{U}$. Besides, it induces the relationship $\eta_j^2-\eta_j=\zeta_j$ between their respective eigenvalues η_j and ζ_j, and we are looking for the positive solution of η_j. To do so, first remark that we may assume that $\boldsymbol{\Omega}_{\mathbf{yx}}(\lambda_2\mathbf{L}+\mathbf{S}_{\mathbf{xx}})\boldsymbol{\Omega}_{\mathbf{xy}}$ and $\mathbf{S}_{\mathbf{yy}}$ are positive definite, when $\boldsymbol{\Omega}_{\mathbf{xy}}\neq\mathbf{0}$ and $n\geq q$; and second, recall that the eigenvalues of a matrix being the product of two positive definite matrices are positive. Hence, $\zeta_j>0$ and the positive solution of η_j is $\eta_j=(1+\sqrt{1+4\zeta_j})/2$. We obtain thus

$$\hat{\boldsymbol{\Omega}}_{\mathbf{yy}}=\mathbf{U}\operatorname{diag}(\boldsymbol{\eta})\mathbf{U}^{-1}\mathbf{S}_{\mathbf{yy}}^{-1}. \tag{17}$$

Direct inversion leads to (12b). To get an expression for $\hat{\boldsymbol{\Omega}}_{\mathbf{yy}}$ which does not require additional matrix inversion, just remark that

$$\mathbf{S}_{\mathbf{yy}}^{-1}=\left(\boldsymbol{\Omega}_{\mathbf{yx}}\hat{\boldsymbol{\Sigma}}_{\mathbf{xx}}^{\lambda_2}\boldsymbol{\Omega}_{\mathbf{xy}}\mathbf{S}_{\mathbf{yy}}\right)^{-1}\boldsymbol{\Omega}_{\mathbf{yx}}\hat{\boldsymbol{\Sigma}}_{\mathbf{xx}}^{\lambda_2}\boldsymbol{\Omega}_{\mathbf{xy}}=\mathbf{U}\operatorname{diag}(\boldsymbol{\zeta}^{-1})\mathbf{U}^{-1}\boldsymbol{\Omega}_{\mathbf{yx}}\hat{\boldsymbol{\Sigma}}_{\mathbf{xx}}^{\lambda_2}\boldsymbol{\Omega}_{\mathbf{xy}}$$

which, combined with (17), leads to (12a).

A.3. Derivation of Proposition 4. To apply the results developed in Tibshirani and Taylor (2012) that rely on the well-known Stein's Lemma (Stein, 1981), we basically need to recast our problem as a classical LASSO applied on a Gaussian linear regression model so as the response vector follows a normal distribution of the form $\mathcal{N}(\boldsymbol{\mu},\sigma\mathbf{I})$. This is straightforward

by means of Proposition 2: in the same way as we derive expression (11), we can go a little farther and reach the following LASSO formulation

$$\arg\min_{\boldsymbol{\omega}\in\mathbb{R}^{pq}} \frac{1}{2}\left\|\begin{pmatrix}\mathbf{A}\\ \sqrt{\lambda_2}\tilde{\mathbf{L}}^{1/2}\end{pmatrix}\boldsymbol{\omega}-\begin{pmatrix}\mathbf{b}\\ \mathbf{0}\end{pmatrix}\right\|_2^2+\lambda_1\|\boldsymbol{\omega}\|_1, \tag{18}$$

with $\mathbf{A}, \mathbf{b}$ and $\tilde{\mathbf{L}}$ defined as in (11). From model (4), it is not difficult to see that $\mathbf{b}$ corresponds to an uncorrelated vector form of $\mathbf{Y}$ so as

$$\begin{pmatrix}\mathbf{b}\\ \mathbf{0}\end{pmatrix}=\begin{pmatrix}-\mathrm{vec}(\mathbf{Y}\boldsymbol{\Omega}_{\mathbf{yy}}^{1/2})\\ \mathbf{0}\end{pmatrix}\sim\mathcal{N}(\boldsymbol{\mu},\mathbf{I}_{nq}).$$

The explicit form of $\boldsymbol{\mu}$ has no interest here. The point is essentially to underline that the response vector is uncorrelated in (18), which allows us to apply Theorem 2 of Tibshirani and Taylor and results therein, notably for the Elastic-Net: by these means, an unbiased estimator of the degrees of freedom in (18) can be written as a function of the active set $\mathcal{A}$ in $\hat{\boldsymbol{\omega}}^{\lambda_1,\lambda_2}$:

$$\hat{\mathrm{df}}_{\lambda_1,\lambda_2}=\mathrm{tr}\left(\mathbf{A}_{\mathcal{A}}\left((\mathbf{A}^T\mathbf{A}+\lambda_2\tilde{\mathbf{L}}\,)_{\mathcal{A}\mathcal{A}}\right)^{-1}\mathbf{A}_{\mathcal{A}}^T\right).$$

Routine simplifications lead to the desired result (14).

A.4. Derivation of Proposition 5. Both results are established by recurrence, the first on indices k, the second on indices ℓ. Consider (5) and indices k for instance: this holds true for $k=1$ by straightforward matrix calculus. Now, assume that $\mathbf{D}^{k,\ell}$ is known, and you are looking for $\mathbf{D}^{k+1,\ell}$ by increasing the alphabet $\mathcal{M}_k$ with one letter, so as lexicographical order still holds true: this can be done by adding successively each letter from $\mathcal{M}_1=\{A,C,G,T\}$ to the left of each word of $\mathcal{M}_k$ (first A to every work in $\mathcal{M}_k$, then C and so on). There are now only two possibilities, which corresponds to each term in the right-hand side of (5):

Consider the first term on the one hand: non null elements in $\mathbf{D}^{k,\ell}$ represent the set of motifs with size k being at distance ℓ to each other. If adding a letter to such motifs, we can split the obtained words in four groups, according to the currently added letter (A, C, G or T). Then, within each of these groups (e.g. the group with words starting by A), all the words shared the very same first letter and thus stay at a distance ℓ from each other: this is just meant by the Kronecker multiplication by $\mathbf{D}^{1,0}$, which builds four blocks corresponding to the four aforementioned groups.

On the other hand, consider the second term: if motifs with size k are at distance $\ell-1$ from each other, what represents $\mathbf{D}^{k,\ell-1}$, then we can add any kind of letter and stay at a distance ℓ for motifs with size $k+1$, what is meant by the Kronecker product with $\mathbf{D}^{1,1}$.

REFERENCES

Bach, F., Jenatton, R., Mairal, J. and Obozinski, G. (2012). Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning* **4** 1–106.

Brown, P. J., Fearn, T. and Vannucci, M. (2001). Bayesian Wavelet Regression on Curves with Applications to a Spectroscopic Calibration Problem. *J. Amer. Statist. Assoc.* **96** 398–408.

Chiquet, J., Grandvalet, Y. and Ambroise, C. (2011). Inferring multiple graphical structures. *Stat. and Comp.* **21** 537–553.

de los Campos, G., Hickey, J. M., Pong-Wong, R., Daetwyler, H. D. and Calus, M. P. L. (2012). Whole Genome Regression and Prediction Methods Applied to Plant and Animal Breeding. *Genetics*.

Efron, B. (2004). The estimation of prediction error: Covariance penalties and cross-validation (with discussion). *J. Amer. Statist. Assoc.* **99** 619–642.

Ferreira, M., Satagopan, J., Yandell, B., Williams, P. and Osborn, T. (1995). Mapping loci controlling vernalization requirement and flowering time in Brassica napus. *Theor. Appl. Genet.* **90** 727–732.

Friedman, J., Hastie, T. and Tibshirani, R. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *J. Stat. Softw.* **33** 1–22.

Gasch, A. P., Spellman, P. T., Kao, C. M., Carmel-Harel, O., Eisen, M. B., Storz, G., Botstein, D. and Brown, P. O. (2000). Genomic expression programs in the response of yeast cells to environmental changes. *Mol. Biol. Cell* **11** 4241–4257.

Hans, C. (2011). Elastic net regression modeling with the orthant normal prior. *J. Amer. Statist. Assoc.* **106** 1383–1393.

Harville, D. H. (1997). *Matrix Algebra from a Statistician's perspective*. Springer.

Hebiri, M. and van De Geer, S. (2011). The smooth-lasso and other l1 + l2 penalized methods. *Electron. J. Stat.* **5** 1184–1226.

Hesterberg, T., Choi, N. M., Meier, L. and Fraley, C. (2008). Least angle and ℓ_1 penalized regression: A review. *Statistics Surveys* **2** 61–93.

Kim, S. and Xing, E. P. (2009). Statistical Estimation of Correlated Genome Associations to a Quantitative Trait Network. *PLoS Genetics* **5**.

Kim, S. and Xing, E. P. (2010). Tree-guided group Lasso for multi-task regression with structured sparsity. In *Proceedings of the 27th International Conference on Machine Learning* 543–550.

Kim, S. J., Koh, K., Boyd, S. and D., G. (2009). ℓ_1 Trend Filtering. *SIAM Review, problems and techniques section* **51** 339–360.

Kole, C., Thorman, C., Karlsson, B., Palta, J., Gaffney, P., Yandell, B. and Osborn, T. (2002). Comparative mapping of loci controlling winter survival and related traits in oilseed Brassica rapa and B. napus. *Mol. Breed.* **1** 201–210.

Lajoie, M., Gascuel, O., Lefort, V. and Brehelin, L. (2012). Computational discovery of regulatory elements in a continuous expression space. *Genome Biology* **13**.

Lee, W. and Liu, Y. (2012). Simultaneous Multiple Response Regression and Inverse Covariance Matrix Estimation via Penalized Gaussian Maximum Likelihood. *J. Multivar. Anal.* **111** 241–255.

Li, X., Panea, C., Wiggins, C. H., Reinke, V. and Leslie, C. (2010). Learning "graphmer" motifs that predict gene expression trajectories in development. *PLoS Comput. Biol.* **6** e1000761.

Mardia, K. V., Kent, J. T. and Bibby, J. M. (1979). *Multivariate Analysis*. Academic Press.

Marin, J. M. and Robert, C. P. (2007). *Bayesian Core: A Practical Approach to Com-*

putational Bayesian Statistics. Springer-Verlag: New-York.

Obozinski, G., Wainwright, M. J. and Jordan, M. I. (2011). Support union recovery in high-dimensional multivariate regression. *Ann. Statist.* **39** 1–47.

Osborne, B. G., Fearn, T., Miller, A. R. and Douglas, S. (1984). Application of near infrared reflectance spectroscopy to compositional analys is of biscuits and biscuit doughs. *J. Sci. Food Agr.* **35** 99–105.

Park, T. and Casella, G. (2008). The Bayesian Lasso. *J. Amer. Statist. Assoc.*

Rothman, A., Levina, E. and Zhu, J. (2010). Sparse multivariate regression with covariance estimation. *J. Comp. Graph. Stat.* **19** 947–962.

Shannon, P. (2013). MotifDb: An Annotated Collection of Protein-DNA Binding Sequence Motifs R package version 1.4.0.

Slawski, M., zu Castell, W. and Tutz, G. (2010). Feature selection guided by structural information. *Ann. Appl. Stat.* **4** 1056–1080.

Sohn, K. A. and Kim, S. (2012). Joint Estimation of Structured Sparsity and Output Structure in Multiple-Output Regression via Inverse-Covariance Regularization. *JMLR* **W&CP** 1081–1089.

Städler, N., Bühlmann, P. and Geer, S. (2010). ℓ_1-Penalization for Mixture Regression Models. *Test* **19** 209–256.

Stein, C. (1981). Estimation of the mean of a multivariate normal distribution. *Ann. Stat.* **9** 1135–1151).

Tibshirani, R. (1996). Regression Shrinkage and Selection Via the Lasso. *J. R. Statist. Soc. B* **58** 267–288.

Tibshirani, R. J. and Taylor, J. (2011). The solution path of the generalized lasso. *Ann. Stat.* **39** 1335–1371.

Tibshirani, R. J. and Taylor, J. (2012). Degrees of Freedom in Lasso Problems. *Electron. J. Stat.*

Tibshirani, R., Saunders, M., Rosset, S., Zhu, J. and Knight, K. (2005). Sparsity and smoothness via the fused lasso. *J. R. Statist. Soc. B* **67** 91–108.

Tseng, P. (2001). Convergence of a block coordinate descent method for nondifferentiable minimization. *J. Optim Theory Appl* **109** 475–494.

Tseng, P. and Yun, S. (2009). A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming* **117** 387–423.

Witten, D. M. and Tibshirani, R. (2009). Covariance-regularized regression and classification for high -dimensional problems. *J. R. Statist. Soc. B* **71** 615–636.

Yin, J. and Li, H. (2011). A sparse conditional Gaussian graphical model for analysis of genetical genomics data. *Ann. Appl. Stat.* **5** 2630–2650.

Yuan, X. T. and Zhang, T. (2014). Partial Gaussian Graphical Model Estimation. *Information Theory, IEEE Transactions on* **60** 1673-1687.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. R. Statist. Soc. B* **67** 301–320.

LaMME - UMR 8071 CNRS/UEVE
Boulevard de France
91000, Évry, FRANCE
E-mail: julien.chiquet@genopole.cnrs.fr
URL: http://stat.genopole.cnrs.fr

UMR 518 AgroParisTech/INRA
rue Claude Bernard
70005 Paris, FRANCE
URL: http://www.agroparistech.fr/mia/